

AD A 125684

## A RAND NOTE

INFERNO: A CAUTIOUS APPROACH  
TO UNCERTAIN INFERENCE

J. R. Quinlan

September 1982

N-1898-RC


DTIC FILE COPY

**Rand**  
SANTA MONICA, CA. 90406

DISTRIBUTION STATEMENT A  
Approved for public release:  
Distribution Unlimited

83 03 14 181

# DISCLAIMER NOTICE



**THIS DOCUMENT IS BEST  
QUALITY AVAILABLE. THE COPY  
FURNISHED TO DTIC CONTAINED  
A SIGNIFICANT NUMBER OF  
PAGES WHICH DO NOT  
REPRODUCE LEGIBLY.**

This publication was supported by The Rand Corporation as part of its program of public service.

The Rand Publications Series: The Report is the principal publication documenting and transmitting Rand's major research findings and final research results. The Rand Note reports other outputs of sponsored research for general distribution. Publications of The Rand Corporation do not necessarily reflect the opinions or policies of the sponsors of Rand research.

Published by The Rand Corporation

**A RAND NOTE**

INFERNO: A CAUTIOUS APPROACH  
TO UNCERTAIN INFERENCE

J. R. Quinlan

September 1982

N-1898-RC

DTIC  
ELECTRIC  
MAR 15 1983  
S H D

**Rand**  
SANTA MONICA, CA. 90406

**DISTRIBUTION STATEMENT A**


Approved for public release;  
Distribution Unlimited

PREFACE

A primary concern of Rand's Information Processing Systems program is the development of Artificial Intelligence technology and its employment in problem domains of practical importance. This research, supported internally by Rand, was undertaken to document the methods that today's expert systems use as they grapple with inexact but valuable knowledge, and to suggest a new approach that avoids some of the evident shortcomings. The Note should be of interest to knowledge engineers who are designing expert systems, and to those users of expert systems worried about how much credence to give to generated results.

SUMMARY

Expert systems commonly employ some means of drawing inferences from domain and problem knowledge, where both the knowledge and its implications are less than certain. Methods used include subjective Bayesian reasoning, measures of belief and disbelief, and the Dempster-Shafer theory of evidence. Analysis of systems based on these methods reveals important deficiencies in areas such as the reliability of deductions and the ability to detect inconsistencies in the knowledge from which deductions were made. A new system called INFERNO addresses some of these points. Its approach is probabilistic but makes no assumptions whatsoever about the joint probability distributions of pieces of knowledge, so the correctness of inferences can be guaranteed. INFERNO informs the user of inconsistencies that may be present in the information presented to it, and can make suggestions about changing the information to make it consistent. An example from a Bayesian system is reworked, and the conclusions reached by that system and INFERNO are compared.



Accession For	
NIIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	<i>As file</i>
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
<i>A</i>	

ACKNOWLEDGMENTS

INFERNO grew out of a study of Frederick Hayes-Roth's WAND, and it incorporates many ideas developed in that system. I gratefully acknowledge the importance of the insights provided by Norman Shapiro and Donald Waterman of Rand, Donald Michie and Tim Niblett of Edinburgh University, and John Reiter and Peter Cheeseman of SRI International.

CONTENTS

PREFACE .....	iii
SUMMARY .....	v
ACKNOWLEDGMENTS .....	vii
FIGURES AND TABLES .....	xi
Section	
I. INTRODUCTION .....	1
II. KEY IDEAS IN INFERENCE NETWORK SYSTEMS .....	4
Uncertain Knowledge About Propositions .....	7
Assumptions .....	10
Control Structure for Propagating Information .....	16
Inconsistent Information .....	17
III. INFERNO .....	21
Representation .....	22
Propagation .....	24
Consistency and Rectification .....	30
IV. A COMPLETE EXAMPLE .....	37
V. CONCLUSION .....	42
Appendix	
A. Derivation of Propagation Constraints .....	47
B. Constraints for Backing Up Inconsistencies .....	50
C. INFERNO Run with AL/X Example .....	52
REFERENCES .....	59



FIGURES

1. An AL/X inference network .....	5
2. A small inference network .....	28
3. INFERNO output .....	29
4. Output with inconsistencies .....	31
5. Changes and rectifications .....	35

TABLES

1. INFERNO relations and their interpretation .....	23
2. INFERNO propagation constraints .....	25
3. AL/X case for analysis .....	39
4. Comparison of findings by AL/X and INFERNO .....	40

## I. INTRODUCTION

The central postulate of knowledge engineering is that systems achieve expert performance from rich, diverse knowledge bases rather than from clever algorithms. The process of solving a problem in some application area is seen as the task of combining domain knowledge and information specific to the problem so that appropriate conclusions can be drawn. As attention continues to shift to real-world problems of practical importance, however, it has become apparent that the corresponding domain and problem knowledge is usually less than certain. Feigenbaum states, "Experience has shown us that [expert] knowledge is largely heuristic knowledge--mostly 'good guesses' and 'good practice' in lieu of facts and rigor" [Feigenbaum, 1979, p. 7]. On the other hand, traditional methods of forming inferences are derived from logic and use techniques such as the Resolution Method [Robinson, 1965] that deal only with categorical information. This Note is concerned with how uncertain knowledge can be used to find inferences that are well-founded even if they are not categorical.

We start with a collection of propositions consisting of facts, expectations, and hypotheses relevant to some problem domain. Subgroups of the propositions are bound together by logical, causal, or other mechanisms, so that knowing something about one proposition may have consequences for others. The information available about the propositions is not clear-cut, but vague, uncertain, or probabilistic. The task of inference under uncertainty is then to maintain an integrated "world view" represented by the sum of deduced and given

information about the propositions as more data are gathered or alternative hypotheses are postulated. Early examples of problem domains analyzed in this way include medical diagnosis and therapy [Shortliffe and Buchanan, 1975], where propositions are symptoms, diseases, and treatments; and geological analysis [Duda, Hart, Nilsson, Reboh, Slocum, and Sutherland, 1977], with propositions about levels of geological structure, geography, and exploitable mineralizations.

A useful model of this type of problem is the inference network [Duda, Gaschnig, and Hart, 1979; Gaschnig, 1981; Hayes-Roth, Waterman, and Lenat, 1978], in which propositions are represented as nodes, each having some measure of validity in the light of information gathered to date. Interdependencies among propositions become links among combinations of nodes. When information is injected at some point in the system, the links allow it to be propagated to all nodes to which it is directly or indirectly relevant and enable the validity measure of these nodes to be altered in the process. When all the propagation ripples have subsided, the new state of the network represents what is known directly about the propositions or can be deduced from their relationships to other propositions.

Section II explores some key ideas relevant to current inference network systems. These lead to the conclusion that another approach is warranted, and Section III introduces a new system called INFERNO. The two major contributions of INFERNO are the guaranteed validity of any inferences that it makes and its concern for, and assistance in establishing, the consistency of the information about the problem and its domain. Section IV illustrates the use of INFERNO and compares it

with one of the more powerful Bayesian systems in common use. Section V summarizes and evaluates INFERNO's contribution to inference under uncertainty and suggests directions for further work.

## II. KEY IDEAS IN INFERENCE NETWORK SYSTEMS

Figure 1 shows a simple inference network adapted from the AL/X model of Reiter (1980) for diagnosing faults that arise on oil-drilling rigs. In addition to providing a good example of many of the important concepts, this particular network is used for a comparative study in Sec. IV. The propositions or nodes are represented by boxes with interrelationships given by directed links. The numbers that appear on top of the boxes are the validity measures for the propositions, in the form of probabilities that the proposition is true; their initial values are the respective prior probabilities. Links are of two kinds: the first kind correspond to antecedent-consequent implications and have two values associated with them. If the antecedent proposition is found to hold, the odds of the consequent proposition are multiplied by the first value; if the antecedent is found to be false, the odds of the consequent are multiplied by the second value; and if the antecedent is found to have some intermediate probability  $p$ , the odds of the consequent proposition are multiplied by an interpolated value determined by  $p$ , the prior probability of the antecedent, and the two values. The second type of link has no values and is used to produce propositions that are Boolean combinations of other propositions. The probability of these nodes is determined by the probabilities of their constituent propositions.

The information in this network is general domain knowledge, in that it is not restricted to the problem of diagnosing a single fault occurrence. Typically, most knowledge codified in an inference network

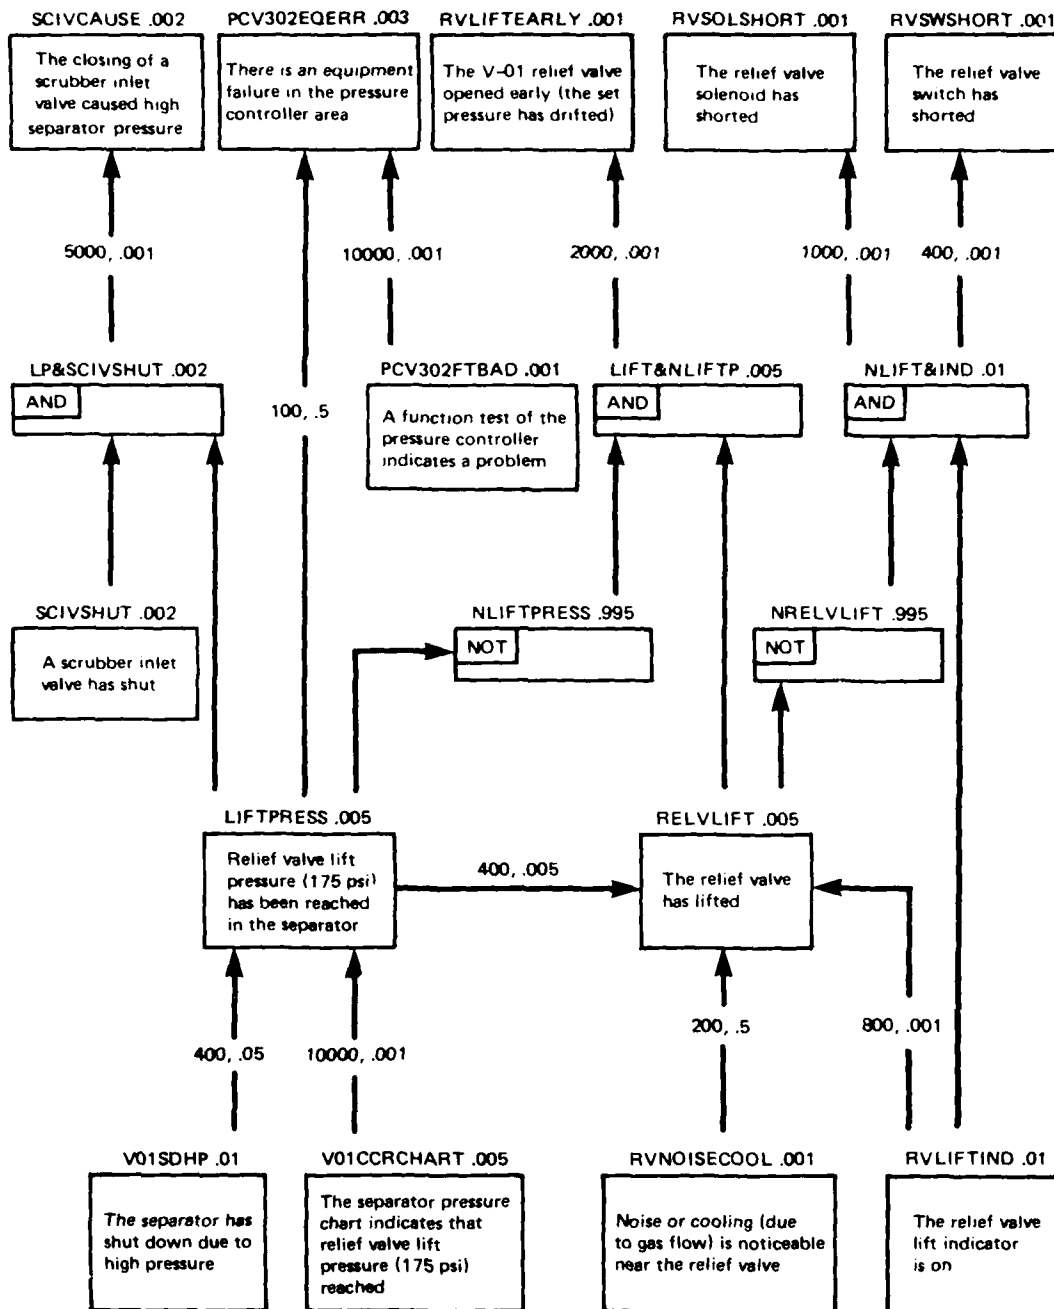


Fig. 1 - An AL/X inference network

is applicable to a range of problems in the domain and usually includes the interrelationships among all propositions of interest. This information is often tedious to set up and debug but can then be used repeatedly. To use the network on a particular problem, the general or domain information is augmented with specific observations, hypotheses, etc., such as "VOISDHP has occurred with probability 0.95." The two bodies of information, specific and general, jointly allow inferences to be drawn for the particular case.

As this example illustrates, the information in an inference network is divided between the part relevant to propositions that is associated with nodes, and the part relevant to relations that resides on the links among nodes. The distinction is important in practice, because typically only the values associated with propositions can be altered as a result of inference. With a few exceptions, such as PI [Friedman, 1981], systems do not allow us to deduce the existence of a connection between nodes.<sup>1</sup> Most systems do not even permit parameters characterizing a relation (such as the values of an antecedent-consequent link) to be modified. This is presumably because the specification of interdependencies belongs to the general information relevant to the domain and so is not expected to change from case to case. The example also shows the two forms of interrelationship that seem to be standard in all systems: expressing one proposition as a

---

<sup>1</sup>If it can be deduced that the proposition  $A \vee B$  is true, this is equivalent to deducing that the relation " $\neg A$  implies  $B$ " holds. However, many systems use a form of one-directional inference (discussed later) that precludes conclusions about  $A$  or  $B$  being drawn from Boolean combinations such as  $A \vee B$ , with the result that the consequences of this sort of ersatz relation are not discovered.

Boolean function of other propositions, and a kind of attenuated implication of the form "if A, then (to extent E) conclude B".

In other respects, inference network systems show considerable dissimilarity. The principal differences can be reduced to four areas of special functional significance: (1) the way in which uncertain information about propositions is represented, (2) the assumptions that form the basis for propagating information, (3) the control structure used for this propagation, and (4) the treatment of inconsistent information. The following discusses alternative approaches in these areas as exemplified in several systems. To avoid continual repetition of references, the principal models used here are Prospector [Duda, Hart, and Nilsson, 1976; Gaschnig, 1980, 1981]; PI [Friedman, 1981]; the system described in Garvey, Lowrance, and Fischler (1981); WAND [Hayes-Roth, unpublished]; SPERIL [Ishizuka and Yao, 1981]; Konolige's information-based approach [Konolige, 1982; Duda, Hart, Konolige, and Reboh, 1979]; AL/X [Reiter, 1980, 1981]; MYCIN [Shortliffe, 1975]; and EMYCIN [van Melle, 1979, 1980]. The widely known work of Zadeh (1979) has been excluded from this list because his "possibilistic" approach using linguistic quantification is so different from the above systems that comparison is difficult.

#### UNCERTAIN KNOWLEDGE ABOUT PROPOSITIONS

There are three common ways of representing uncertain information about propositions. The most straightforward approach (employed in Prospector and AL/X and used in the sample network) provides a



subjective probability for each proposition.<sup>2</sup> If E is the sum of all information available to the system, each proposition A has associated with it a value representing the probability  $P(A|E)$  that A holds in the light of E, derived under some regimen such as Bayesian inference. The main point here is not how this value is derived, but the fact that what we know or have inferred about a proposition is represented by a single value.

Two criticisms have been advanced concerning this approach. The first is that the single value tells us nothing about its precision, which may be very low when the value is derived from uncertain evidence. To say that  $P(A|E)$  is 0.5 might mean either  $0.500 \pm 0.001$  or  $0.5 \pm 0.3$ , two very different pieces of information. Barnett writes, "There is trouble representing what we actually know without being forced to overcommit when we are ignorant" [Barnett, 1981]. Garvey et al. note, "A likelihood represented by a point probability value is usually an overstatement of what is actually known" [Garvey et al., 1981]. The second criticism is that the single value combines the evidence for and against A without indicating how much there is of each. Intuitively, the probability of A in the light of E might have the same value when no evidence in E is relevant to A as when E contains strong but counterbalancing arguments for A and against A.

MYCIN uses an alternative approach in which there are two separate values for the validity of each proposition.  $MB[A,E]$  is a probability-like measure of belief in A given E, and  $MD[A,E]$  is a similar value for

---

<sup>2</sup>The adjective "subjective" means that the probability does not have to come from ideas of relative frequency in the limit but can reflect fair odds on the proposition being true.

disbelief in A given E. If E could be partitioned into two parts,  $E^+$  favoring A and  $E^-$  opposing A,  $MB[A,E]$  corresponds in intent with  $P(A|E^+)$  and  $MD[A,E]$  with  $P(\sim A|E^-)$ . However, the belief and disbelief measures are independent and so cannot be probabilities, although they have the same interpretation at their extremes; if  $MB[A,E]$  is 1, E provides incontrovertible evidence that A is true. The two measures are combined into a single assessment of A in the light of E, called the certainty factor  $CF[A,E]$  and defined as  $MB[A,E] - MD[A,E]$ .

The two-value approach is also subject to the criticism about precision, since both of the belief measures are point values. It does, however, overcome the second objection because the interplay of evidence pro and con is manifest. This separation can be nullified if the belief and disbelief measures are used not as distinct entities but only as an amalgamated certainty (as is the case in MYCIN's successor, EMYCIN), because the amalgamation restores what is essentially the original single-value system. The scheme suffers from a new disability, however, in that there is no foundation of theory underpinning and justifying the interpretation and weighing of separate belief and disbelief measures.

The third approach is used in systems with different pedigrees, including those that employ the Dempster-Shafer theory of evidence (such as Garvey et al., 1981; Barnett, 1981) and others (SPERIL, WAND). Instead of representing the probability of a proposition A by a point value, this approach bounds the probability to a subinterval  $[s(A), p(A)]$  of  $[0,1]$ . The exact probability  $P(A)$  of A may be unknown but bounded by  $s(A) \leq P(A) \leq p(A)$ . The precision of our knowledge about A is immediately plain, with our uncertainty characterized by the difference

$p(A) - s(A)$ . If this is small, our knowledge about A is relatively precise; if it is large, we know correspondingly little. If  $p(A)$  equals  $s(A)$ , our knowledge about A is exact and reverts to the point probability of the first approach. Notice that the inequality above can be recast as two assertions: (1) that the probability of A is at least  $s(A)$ , and (2) that the probability of  $\sim A$  is at least  $1 - p(A)$ . Thus this representation also addresses the second criticism, because it keeps what amount to separate measures of belief and disbelief in A derived from the available evidence. Finally, as with the first scheme, there is the solid ground of probability theory on which to base the interpretation of the values  $s(A)$  and  $p(A)$ .

#### ASSUMPTIONS

We now turn to some of the assumptions underlying the ways in which information is propagated in the network. Recent findings have cast doubts on the appropriateness of many of these assumptions, and thus indirectly on the methods of propagation that use them. The position from which these assumptions are assessed is a conservative one: rather than rely on questionable assumptions, it is preferable to accept the penalty of less definitive inferences.

The systems that use Bayes' theorem provide a good starting point. A seminal paper of Duda, Hart, and Nilsson (1976) is meticulous in setting out explicitly the basis for the propagation scheme used in Prospector and subsequently in AL/X. One trouble arises when two distinct pieces of evidence  $E_1$  and  $E_2$  are relevant to a proposition A. In order to update our assessment of A we need to compute  $P(A|E_1 \& E_2)$ ,

but knowing how to update  $A$  separately for each of  $E_1$  and  $E_2$  such as by knowing each of  $P(E_1)$ ,  $P(E_2)$ ,  $P(A \& E_1)$ , and  $P(A \& E_2)$  is not sufficient to determine this. The general case would require propagation parameters (such as the upper and lower values on links in the earlier example) for every subset of inputs to a proposition or, equivalently, the complete joint distribution of all propositions. In many large inference networks, asking the system designer to specify separate values for each possible combination of evidence relevant to a proposition would overtax his knowledge and presumably his patience, although some systems such as PIP [Szolovits and Pauker, 1978] do require just this. Prospector and AL/X sidestep this problem by making the conditional independence assumptions

$$P(E_1 \& E_2 | A) = P(E_1 | A) \times P(E_2 | A)$$

$$P(E_1 \& E_2 | \sim A) = P(E_1 | \sim A) \times P(E_2 | \sim A)$$

and  $P(A | E_1 \& E_2)$  can now be determined under these constraints. Although they seem reasonable at first sight, Szolovits and Pauker report that "The assumption of conditional independence is usually false" [Szolovits et al., 1978, p. 121]. Pednault, Zucker, and Muresan go further still [Pednault et al., 1981]. In consultation systems, the case-specific information can be often regarded as evidence on which "to distinguish among competing hypotheses" [Duda et al., 1976, p. 1], where the hypotheses are a subset of the propositions. In particular, if there are three or more mutually exclusive and exhaustive propositions to which the evidence is apparently relevant, Pednault et al. prove that conditional independence implies strict independence. They then cite a

theorem proving that conditional independence plus strict independence is sufficient to establish that the evidence is really irrelevant to the propositions! Since this is a preposterous conclusion, presuming the network builder knew his onions, the contradiction shows that the conditional independence assumption must have been false.<sup>3</sup>


Systems of the MYCIN family and those based on the Dempster-Shafer theory of evidence both make similar assumptions for the purpose of combining evidence. The former employ formulas that mirror the calculation of probabilities of independent events; for example,

$$MB[A,E1\&E2] = MB[A,E1] + MB[A,E2] - MB[A,E1] \times MB[A,E2]$$

which seems to indicate that  $MB[A,E1]$  and  $MB[A,E2]$  are taken to be independent. The latter uses a combination rule called the "orthogonal sum" that assumes the evidence being combined is independent even if it is imprecise [Barnett, 1981]. But we observe in general that unless the inference network resembles a tree, there will be one or more cases where  $E1$  and  $E2$  are relevant to  $A$  but the sets of propositions indirectly relevant to  $E1$  and  $E2$  are not disjoint. Diagrammatically, this is

---

<sup>3</sup>A recent paper [Pearl, 1982] shows how this difficulty can be circumvented if propositions can be generalized to multivalued variables. If  $A$ ,  $B$ , and  $C$  are mutually exclusive and complete propositions, they can be replaced by a single proposition  $H$  with values  $A$ ,  $B$ , and  $C$ . The corresponding conditional independence assumption is thereby weakened, and independence of the pieces of evidence relevant to  $H$  is no longer implied.



Situations such as this are the rule rather than the exception in practical inference networks--see for example those given by Gaschnig (1980)--and in such cases, E1 and E2 are clearly not independent. Since once again the propagation schemes depend on an assumption that is at least questionable, the inferences reached via such propagation also contain seeds of doubt.

Konolige has developed a novel approach to Bayesian inference that does not require the conditional independence assumption. Imagine a system with only three propositions, A, B, and C. The joint distribution of these propositions is the set of eight elementary probabilities  $P(A \& B \& C)$ ,  $P(A \& B \& \sim C)$ ,  $P(A \& \sim B \& C)$ , ...  $P(\sim A \& \sim B \& \sim C)$ . The usual sorts of information specified in inference nets, such as prior and conditional probabilities  $P(A)$  and  $P(A|B)$ , can be mapped into corresponding linear constraints on these elementary probabilities. In the absence of assumptions such as conditional independence, the joint distribution is underconstrained, and we will let  $\{Z\}$  denote the set of

distributions satisfying the constraints. Choice of different joint distributions from  $\{Z\}$  will in general assign different posterior probabilities to the propositions. Konolige argues that the best choice of a single distribution from the candidates in  $\{Z\}$  is the one that contains the least additional information about dependencies among the propositions, and he shows that this is equivalent to selecting the candidate with maximum entropy. While the approach is clearly a powerful and interesting one, difficulties remain. First, the best choice of a candidate distribution may still happen to be incorrect and thereby misleading, especially if the range of candidates in  $\{Z\}$  is large. Second, the results can be sensitive to the way in which propositions are formulated. As a trivial illustration, a single unconstrained proposition will be assigned a probability of 0.5, but if it is expressed as the conjunction of two unconstrained propositions, its probability will be taken as 0.25. Third, the necessary computations are feasible only if the problem can be decomposed into small overlapping groups of propositions and if all constraints on the joint distribution are linear (e.g., propositions cannot be asserted to be independent).

Another type of assumption concerns the way Boolean combinations of propositions are handled. Prospector, AL/X, and many others use the "fuzzy" formulas for conjunction and disjunction:

$$P(A \vee B | E) = \max( P(A | E), P(B | E) )$$

$$P(A \& B | E) = \min( P(A | E), P(B | E) )$$

This gives the most pessimistic estimate possible for the probability of

the disjunction but the most optimistic estimate for the probability of the conjunction. It is unclear why a consistent approach should not be taken, for example by using the corresponding optimistic estimate for disjunction:

$$P(A \vee B | E) = \min(1, P(A | E) + P(B | E))$$

The MYCIN and PI computations of belief and disbelief for conjunctions and disjunctions take a similar form to the fuzzy formulas, so the same criticism applies.

Finally, there is the question of how to update a consequence of proposition A when A is known with less than certainty. The approach taken in both the MYCIN and Prospector families is to interpolate from the case where A is true.<sup>4</sup> Several interpolation schemes are discussed by Duda et al. (1976) and more recently by Paterson (1981). Choice of a particular scheme seems to be a matter of taste, and it is unclear in practice what effects the different schemes have on the conclusions reached through chains of inferences.

In summary, current systems typically depend on assumptions of one form or another. If these assumptions turn out to be unjustified in a particular application, then the inferences drawn in that case are erroneous to some degree. On the other hand, no system as yet seems to allow the user to provide information about the independence of subsets of the propositions when that independence is known. For example, if propositions A and B are known to be independent, it would seem beneficial to be able to assert (as opposed to assume) this information;

---

<sup>4</sup>Konolige's approach manages to avoid this difficulty also.




the probability of A&B could be then computed accurately rather than from either optimistic or pessimistic estimates.

#### CONTROL STRUCTURE FOR PROPAGATING INFORMATION

Many inference network systems were designed to operate in a consulting environment. They therefore draw a clear distinction between nodes representing propositions about which the system must be informed and those ("hypotheses") that depend on other propositions. Such consultation systems do not allow information in a case to be volunteered about a hypothesis; the user can provide input concerning only those propositions that cannot be inferred from others. The distinction is significant because the flow of inference is constrained to a single direction, from the propositions that constitute the "raw evidence" to the furthest removed propositions (often called "goals"). The AL/X example of Fig. 1 illustrates this approach. Data can be given concerning only the six propositions that have no input from other propositions, and the inferences from this information follow the directions of the links. Thus, like most systems, AL/X does not allow a hypothesis to be posited for the purpose of drawing conclusions about the pattern of evidence that might be expected to support it.

The distinction is carried one step further in Garvey et al. (1981), where the "raw evidence" is regarded as coming from sensors. Internal nodes whose only input comes from sensor nodes are treated differently from other nodes, as illustrated below:



When two or more sensor nodes provide evidence supporting hypothesis A, this evidence is combined using Dempster's orthogonal sum mentioned previously; S1 and S2 are taken as independent. But when two internal nodes such as A and B provide (inferred) evidence supporting Z, no such assumption is made, and all but the strongest evidence is ignored.

#### INCONSISTENT INFORMATION

The idea of evidence for and against a proposition was introduced in the discussion of how information is represented. One particular case arises when both sets of evidence are fairly convincing, i.e., where the given information supports powerful arguments that proposition A is true but also counterarguments showing that A is false. Since both the arguments and counterarguments are derived from the same given information, the conflict is implicit in the information itself, and the information can well be labeled inconsistent. Situations of this kind arise frequently in expert domains--after all, most specialist areas exhibit well-known controversies--so they should also be anticipated in expert systems for such domains. Yet an examination of current inference network systems reveals significant deficiencies in detecting inconsistency and in treating it appropriately.

The conflict of evidence pro and con cannot even be detected in the straightforward Bayesian approach typified by Prospector and AL/X. As was illustrated in Fig. 1, the input for a particular problem consists of certainties or probabilities for some or all of the input propositions. These certainties are used to determine multipliers for the odds of higher hypothesis-type propositions, and thus information propagates. Any combination of values for the input certainties is as acceptable as any other as far as the propagation machinery is concerned; the only effect is to change some or all of the computed certainties for the hypotheses. As a result, the concept of conflicting problem information does not exist in this formalism! The position regarding the domain information, however, is not so simple. Specifying the prior probability of each proposition and two odds multipliers for each inferential link represents an overconstraint of the system [Duda et al., 1976], so there would seem to be a mechanism for detecting conflict. The trouble is that inconsistency of this type (produced by overconstraint) is almost inevitable, so consistency is not an achievable goal. In fact, the piecewise linear interpolation function used by Prospector can be regarded as a way of living with inconsistency. The case is clearer still when we turn to Konolige's modified system (discussed previously). Here, any nondegenerate system is guaranteed to be underconstrained, so inconsistency of problem or domain information is simply not possible.

Recall that MYCIN keeps separate measures of belief  $MB[A,E]$  and disbelief  $MD[A,E]$  for every proposition  $A$  in the light of evidence  $E$ . On the surface, this offers the possibility of detecting inconsistency,

but since MYCIN's measures are not probabilities there is no indication of when these values are irreconcilable. For example, can the measures be said to be in conflict when they both have a value of 0.6? Only when both measures are 1 is there a proven contradiction, and in this unlikely event, MYCIN would accept only the first finding and would ignore evidence implying the second. MYCIN's descendent EMYCIN has abandoned the two-valued approach and so would be unable to detect any contradiction.

It is only in systems using a two-valued approach where the values are probabilities that there is a firm basis for detecting general inconsistency. If the probability of proposition A lies in the interval  $[s(A), p(A)]$  and  $s(A)$  is greater than  $p(A)$ , the lines of reasoning leading to these bounds clearly make use of inconsistent information in the form of relations among propositions and/or assertions about the probability of propositions. Inconsistency can be detected at different stages. For example, Earnett (1981) and Garvey et al. (1981) set out conditions that must be satisfied if two pieces of evidence relevant to a hypothesis are consistent. WAND accepts information incrementally and detects that a datum is inconsistent with previous input when it would lead to unsatisfiable bounds of the type described above.

Even more important than detecting inconsistency is the question of what to do about it when it occurs. No system for uncertain inference seems to have arrived at a satisfactory answer. For example, the Dempster-Shafer systems cited above do not appear to address the question, and in fact the orthogonal sum operation that they use to combine evidence breaks down if the evidence is inconsistent. Again,

WAND will simply refuse to accept a datum found to contradict previous data, even though the previous information may be the cause of the problem.

### III. INFERNO

There were several motivations for producing a new system, all of which are apparent from the previous section:

1. An inference system should not depend on any assumptions about the probability distributions of the propositions.
2. Conversely, it should be possible to assert common relationships between propositions (such as independence) when the relationships are indeed known.
3. There should be no distinction between propositions; it should be possible to posit information about any set of propositions and observe the consequences for the system as a whole.
4. If the information provided to the system is inconsistent, this fact should be made evident along with some notion of alternative ways that the information could be made consistent.

The first point is obvious enough: if no assumptions are made, no errors will be generated (and propagated) by the process of deriving inferences. This does not prevent the drawing of erroneous conclusions from faulty information, but it does guarantee that any errors that arise will be attributable to the data and not to the system. In the absence of assumptions, the inferences that are made may be weak, and requirement (2) above enables them to be strengthened where it is safe to do so. Requirement (3) concerns the ways that the system might be used. It should be possible both to reason backward from hypothetical situations and to reason forward from observations to conclusions using

the same proposition interrelationships. Consequently, it must be possible to propagate information in all directions from a node, not just towards a goal. As a simple example, from knowing A and B the system must be able to deduce  $A \& B$ ; from knowing  $A \vee B$  and  $\sim A$  the system must be able to deduce B. The final requirement is that the system must be aware of the presence of inconsistent information but be able to accept and propagate it nonetheless. INFERNO flags propositions about which inconsistent inferences can be drawn and incorporates a mechanism akin to dependency-directed backtracking [Stallman and Sussman, 1977] for suggesting changes to the data that are sufficient to remove all inconsistencies.

#### REPRESENTATION

The decision that no assumptions are to be used when propagating information immediately rules out using point probabilities in the general case. INFERNO uses a two-value scheme similar in intent to the interval approach  $[s(A), p(A)]$  above. It turns out that INFERNO can more easily be described if the two values characterizing a proposition A are  $t(A)$  and  $f(A)$ , where

$$P(A) \geq t(A) \quad \text{and} \quad P(\sim A) \geq f(A)$$

i.e.,  $t(A)$  is a lower bound on the probability of A derived from the evidence for A, and  $f(A)$  is a lower bound on  $\sim A$  derived from the evidence against A. Evidence is for A if it allows the inference that  $P(A) \geq X$  and against A if it gives  $P(A) \leq X$ . We define (the information about) proposition A to be consistent as long as

$$t(A) + f(A) \leq 1$$

in which case,  $s(A) = t(A)$  and  $p(A) = 1 - f(A)$ .

INFERNO uses relations among propositions that are patterned on and extend those in WAND, although the interpretation of these relations and the ways in which they are used in propagation differ from those in WAND. The relations themselves and their interpretation are given in Table 1; despite their somewhat arbitrary appearance, they seem sufficient to express common interdependencies. In addition to variants of weak implication and Boolean combinations of propositions, the relations permit assertions that sets of propositions are independent or mutually exclusive as in requirement (2) above, although INFERNO currently uses this information only in the context of the relation in which it appears. Notice that inhibits, requires, and unless can be

Table 1

INFERNO RELATIONS AND THEIR INTERPRETATION

<u>Relation</u>	<u>Interpretation</u>
A enables S with strength X	$P(S A) \geq X$
A inhibits S with strength X	$P(\sim S A) \geq X$
A requires S with strength X	$P(\sim A \sim S) \geq X$
A unless S with strength X	$P(A \sim S) \geq X$
A negates S	$A \equiv \sim S$
A conjoins $\{S_1, S_2, \dots, S_n\}$	$A \equiv \&_i S_i$
A conjoins-independent $\{S_1, S_2, \dots, S_n\}$	$A \equiv \&_i S_i$ ; and for all $i \neq j$ , $P(S_i \& S_j) = P(S_i) \times P(S_j)$
A disjoins $\{S_1, S_2, \dots, S_n\}$	$A \equiv \vee_i S_i$
A disjoins-independent $\{S_1, S_2, \dots, S_n\}$	$A \equiv \vee_i S_i$ ; and for all $i \neq j$ , $P(S_i \& S_j) = P(S_i) \times P(S_j)$
A disjoins-exclusive $\{S_1, S_2, \dots, S_n\}$	$A \equiv \vee_i S_i$ ; and for all $i \neq j$ , $P(S_i \& S_j) = 0$
$\{S_1, S_2, \dots, S_n\}$ mutually exclusive	for all $i \neq j$ , $P(S_i \& S_j) = 0$



defined in terms of enables and negates; the discussion of propagation and rectification will be simplified by lumping the four of them together as "enables-type" relations.

#### PROPAGATION

Each proposition A initially has the trivial bounds  $t(A) = 0$  and  $f(A) = 0$ . These bounds may be changed by explicit information from the case being studied or by inferences from other bounds through the relations connecting A with other propositions. As more information is provided or inferred, the range within which the probability  $P(A)$  of proposition A is known to lie can only become smaller. This is reflected in larger values for one or both of  $f(A)$  and  $t(A)$ .

Suppose that we have just computed a higher value for one of the bounds of a proposition. This new information is propagated by checking all relations in which the proposition is involved and perhaps increasing the bounds of other propositions to ensure that relevant propagation constraints are satisfied. These constraints are summarized relation by relation in Table 2 and are formally derived in Appendix A. They make no assumptions whatsoever about the probability distributions of any propositions and follow mainly from the universal inequality

$$\max P(S_i) \leq P(S_1 \vee S_2 \vee \dots \vee S_n) \leq \sum_i P(S_i)$$

in various guises. Each inequality asserts that some bound is greater than or equal to some expression involving other bounds, and is interpreted as a form of production rule:

Table 2

INFERNO PROPAGATION CONSTRAINTS

A enables S with strength X:

$$t(S) \geq t(A) \times X \quad (1.1)$$

$$f(A) \geq 1 - (1 - f(S)) / X \quad (1.2)$$

A negates S:

$$t(A) = f(S) \quad (2.1)$$

$$f(A) = t(S) \quad (2.2)$$

A conjoins  $\{S_1, S_2, \dots, S_n\}$ :

$$t(A) \geq 1 - \sum_i (1 - t(S_i)) \quad (3.1.1)$$

$$f(A) \geq f(S_i) \quad (3.1.2)$$

$$t(S_i) \geq t(A) \quad (3.1.3)$$

$$f(S_i) \geq f(A) - \sum_{j \neq i} (1 - t(S_j)) \quad (3.1.4)$$

A conjoins-independent  $\{S_1, S_2, \dots, S_n\}$ :

$$t(A) \geq \prod_i t(S_i) \quad (3.2.1)$$

$$f(A) \geq 1 - \prod_i (1 - f(S_i)) \quad (3.2.2)$$

$$t(S_i) \geq t(A) / \prod_{j \neq i} (1 - f(S_j)) \quad (3.2.3)$$

$$f(S_i) \geq 1 - (1 - f(A)) / \prod_{j \neq i} t(S_j) \quad (3.2.4)$$

A disjoins  $\{S_1, S_2, \dots, S_n\}$ :

$$t(A) \geq t(S_i) \quad (4.1.1)$$

$$f(A) \geq 1 - \sum_i (1 - f(S_i)) \quad (4.1.2)$$

$$t(S_i) \geq t(A) - \sum_{j \neq i} (1 - f(S_j)) \quad (4.1.3)$$

$$f(S_i) \geq f(A) \quad (4.1.4)$$

A disjoins-independent  $\{S_1, S_2, \dots, S_n\}$ :

$$t(A) \geq 1 - \prod_i (1 - t(S_i)) \quad (4.2.1)$$

$$f(A) \geq \prod_i f(S_i) \quad (4.2.2)$$

$$t(S_i) \geq 1 - (1 - t(A)) / \prod_{j \neq i} f(S_j) \quad (4.2.3)$$

$$f(S_i) \geq f(A) / \prod_{j \neq i} (1 - t(S_j)) \quad (4.2.4)$$

A disjoins-exclusive  $\{S_1, S_2, \dots, S_n\}$ :

$$t(A) \geq \sum_i t(S_i) \quad (4.3.1)$$

$$f(A) \geq 1 - \sum_i (1 - f(S_i)) \quad (4.3.2)$$

$$t(S_i) \geq t(A) - \sum_{j \neq i} (1 - f(S_j)) \quad (4.3.3)$$

$$f(S_i) \geq f(A) + \sum_{j \neq i} t(S_j) \quad (4.3.4)$$

$\{S_1, S_2, \dots, S_n\}$  mutually exclusive:

$$f(S_i) \geq \sum_{j \neq i} t(S_j) \quad (5.1)$$

if the previous value of the bound on the left-hand side  
is less than the value of the right-hand side  
then the bound is increased to this new value.

Each constraint is activated whenever any bound mentioned in its right-hand side is changed. As an illustration, consider the hypothetical relation "A conjoins {Q,R}." From constraints (3.1.1) through (3.1.4) in Table 2,

- $t(A)$  must be checked whenever  $t(Q)$  or  $t(R)$  is increased.
- $f(A)$  must be checked whenever  $f(Q)$  or  $f(R)$  is increased.
- $t(Q)$  and  $t(R)$  must be checked whenever  $t(A)$  is increased.
- $f(Q)$  must be checked whenever  $f(A)$  or  $t(R)$  is increased, and likewise  $f(R)$  for  $f(A)$  or  $t(Q)$

If one of these bounds is increased in line with the constraints, that increase must be propagated in turn. If the current value of a bound satisfies the constraint, then of course no change need take place and no propagation is involved.

There are two detrimental comments that must be made about this propagation mechanism. First, while the constraints can be derived from the interpretations of Table 1, the converse is not true; the constraints are weaker than the interpretations. Consider, for example,<sup>5</sup> the two relations

A enables B with strength X

C conjoins-independent {A,B}

---

<sup>5</sup>This example is due to Norman Shapiro.

whose interpretation from Table 1 consists of the assertions

$$P(B|A) \geq X$$

$$C \equiv A \& B$$

$$P(A \& B) = P(A) \times P(B)$$

From the latter, one can deduce (inter alia) that  $P(B) \geq X$ . The constraints in Table 2 do not allow this inference because knowledge of the independence of A and B is confined to the single relation that asserts it. The values of the bounds inferred by the propagation mechanism will always be correct but in some cases may be weaker than those that can be derived using the interpretations rather than the constraints.

The second point concerns the termination of propagation. It turns out that if the information about propositions is not consistent, a propagation chain with ever-increasing bounds can arise. Consider the relation "A conjoins-independent {Q,R}" and suppose the bound  $t(Q)$  is increased to a value, X say. By propagation constraint (3.2.1),  $t(A)$  must be increased to  $X \times t(R)$ , and by substituting this value in constraint (3.2.3),

$$t(Q) \geq X \times t(R) / (1 - f(R))$$

If  $t(R) + f(R) > 1$ , this would require  $t(Q)$  to be greater than X, so there is a positive feedback loop. Shapiro provided an elegant demonstration that the same problem can arise even when the data are consistent. It is possible to construct a set of relations whose corresponding constraints are satisfied only by a unique assignment of

irrational values to the bounds. But since the bounds are initially zero and the propagation constraints compute only rational functions of the set of bounds, this unique solution cannot be found by a finite number of such computations. In practice, however, inference network systems usually prohibit changes propagating back to the source of the original disturbance.<sup>6</sup> With this prohibition, propagation will always terminate, even when one or more propositions have inconsistent bounds. The result again is that the bounds computed by INFERNO are correct consequences of the data but may not be the tightest bounds that can be inferred from the data.

The propagation mechanism can be illustrated using the small network shown in Fig. 2. There are five propositions labeled A through E and three relations among them: C is the conjunction of A and B, where nothing is known about the joint distribution of A and B; E is the disjunction of C and D and these are known to be independent; and B directly suggests E with  $P(E|B)$  at least 0.8. Suppose that in some case




Fig. 2--A small inference network

<sup>6</sup>Friedman (1981), for example, describes this prohibition as "an axiom of plausible inference." An undesirable side effect of the prohibition is that any relaxation-style computation of values is precluded, leading to bounds that are weaker than the relations would otherwise support.

$P(C)$  is found to lie in the interval  $\{0.55, 0.65\}$ , i.e.,  $t(C) = 0.55$  and  $f(C) = 0.35$ . The relation "C conjoins {A,B}," constraint (3.1.3), and  $t(C)$  combine to give the inference that  $t(A)$  and  $t(B)$  must be at least 0.55. Since  $t(D)$  and  $f(D)$  are initially zero, the relation "E disjoins-independent {C,D}" and  $t(C)$  give  $t(E) = 0.55$  by constraint (4.2.1). If it is now learned that the probability of B is 0.9 (i.e.,  $t(B) = 0.9$  and  $f(B) = 0.1$ ), the following inferences can be made:

$f(A) = 0.25$  from constraint (3.1.4),  $f(C)$ , and  $t(B)$   
 $t(E) = 0.72$  from constraint (1.1) and  $t(B)$   
 $t(D) = 0.2$  from constraint (4.2.3),  $t(E)$ , and  $f(C)$

The sum of our knowledge of the propositions, both given and inferred, is shown in the fragment of INFERNO output in Fig. 3. Notice that a

---

"A": range 0.55 - 0.75  
LB from  $\text{Pr}("C") \geq 0.55$ , by conjoins  
UB from  $\text{Pr}("B") \geq 0.9$ , by conjoins

"B": range 0.9 - 0.9  
LB by assumption  
UB by assumption

"C": range 0.55 - 0.65  
LB by assumption  
UB by assumption

"D": range 0.2 - 1  
LB from  $\text{Pr}("E") \geq 0.72$ , by disjoins-independent

"E": range 0.72 - 1  
LB from  $\text{Pr}("B") \geq 0.9$ , by enables

---

Fig. 3--INFERNO output

justification is given for each nontrivial lower bound (LB) or upper bound (UB) on the probability of a proposition. This justification is either "by assumption" when the information was supplied to INFERNO or another bound and a relation that together triggered the inference.

#### CONSISTENCY AND RECTIFICATION

When  $t(A) + f(A) > 1$  for some proposition A, the information about A is inconsistent and one or both of the bounds must be incorrect. Since the propagation constraints are provably correct, this inconsistency can arise only from contradictions implicit in the information given to the system, i.e., the interdependencies among the propositions expressed by the relations or the explicit bounds on the probabilities of propositions. Whenever INFERNO detects that a proposition is inconsistent, it places a flag \*\* in front of it. In addition, those constraints whose derivation requires the consistency axiom are weakened by using looser bounds  $t'(A)$  and  $f'(A)$ ;  $t'(A)$  is the smaller of  $t(A)$  and  $1 - f(A)$ , and  $f'(A)$  is defined similarly.

In the example of Fig. 3, it was inferred that  $P(E)$  lay in the interval  $[0.72, 1]$ . What would happen if the system was explicitly informed that  $P(E) \leq 0.5$ ? The INFERNO output in Fig. 4 shows the consequences of this additional assertion. Taking proposition B as an example, we are told that  $t(B) \approx 0.9$ . We can infer from  $f(E)$  and constraint (1.2) that  $f(B)$  is at least 0.375. These bounds on the probability of B are clearly incompatible.

INFERNO can be asked to suggest combinations of changes that will make all the information consistent. Each change takes the form of

---

"A": range 0.55 - 0.6  
LB from Pr("C")  $\geq$  0.55, by conjoins  
UB from Pr("C")  $\leq$  0.5, by conjoins

\*\* "B": range 0.9 - 0.625  
LB by assumption  
UB from Pr("E")  $\leq$  0.5, by enables-inverse

\*\* "C": range 0.55 - 0.5  
LB by assumption  
UB from Pr("E")  $\leq$  0.5, by disjoins-independent

\*\* "D": range 0.2 - 0  
LB from Pr("E")  $\geq$  0.72, by disjoins-independent  
UB from Pr("C")  $\leq$  0.5, by disjoins-independent

\*\* "E": range 0.72 - 0.5  
LB from Pr("B")  $\geq$  0.9, by enables  
UB by assumption

---

Fig. 4--Output with inconsistencies

lowering the value of an externally specified bound or reducing the strength of an enables-type relation. A combination of changes that is sufficient to make the bounds on all propositions consistent is called a rectification.

The first step is to examine each inconsistent proposition A in turn, to see how its bounds could be made consistent. This can be achieved by lowering  $t(A)$  to  $1 - f(A)$ , lowering  $f(A)$  to  $1 - t(A)$ , or changing both bounds to intermediate values; INFERNO currently considers only the first two alternatives. Suppose that we are trying to lower some bound to a value V. There are two possibilities:

- The bound was supplied explicitly as input to the system. In this case, it can clearly be changed only with the consent of the user.



- The bound was inferred from a propagation constraint and the values of one or more antecedent bounds. Once more there are two cases:
  - If the propagation constraint came from an enables-type relation, then the relation itself can be weakened. INFERNO calculates the reduced strength of the relation that would have produced the desired value V.
  - Whatever the nature of the relation, the inferred value is a function of the antecedent bounds. INFERNO looks for lower values of these bounds that would have allowed the lower value V to be inferred.

In the last case, the lower values of the antecedent bounds must then be analyzed in a similar fashion, so that finding changes is a lot like propagation in reverse. The backing-up constraints are similar, but not identical, to the propagation constraints. Appendix B lists them without proof, since their derivation parallels that of the propagation constraints.

The process of finding changes is not exact in the current implementation of INFERNO. First, the system considers only single changes that will make a proposition consistent; trying to reduce the value of  $t(A \vee B)$ , for example, can present a problem. Second, it does not verify that there are no relations other than the one referenced in its justification that also constrain the value of the bound to something greater than V. There is no fundamental difficulty (other than program complexity) in removing these limitations.

The second step involves assembling the changes discovered by this process into rectifications. Each rectification contains one change

from each of the inconsistent propositions, so the conjunction of the changes is sufficient to fix the entire collection of propositions (subject to the caveat above). Some of the inconsistent propositions may have matching or compatible changes associated with them, so there are often fewer changes in a rectification than there are inconsistent propositions.

The third step is to rank the rectifications in order of potential utility. The rule of thumb used here is that small adjustments of values do less violence to the input information than gross alterations and are therefore more likely to be acceptable to the user. As a simple model, we define the reluctance of (the user to accept) a change to be the magnitude of the modification that it entails, either the numeric reduction in strength of an enables-type relation or the difference in the old and new values of a bound. The reluctance of a rectification is similarly taken as the sum of the reluctances of the changes it contains. Alternative rectifications are presented in order of increasing reluctance, so that those suggested first are more likely to be reasonable.

If there are many inconsistent propositions, each with many alternative changes, the process of finding and ranking all possible rectifications is combinatorially explosive, and there would be too many rectifications to display to the user anyway. This problem is circumvented by establishing a parameter ( $R$ , say) so that only the  $R$  best rectifications are ever displayed. The processes of generation and ranking are then combined; if a partial rectification has a reluctance greater than that of all the  $R$  best rectifications found so far, then no

rectification containing the partial rectification need ever be generated. With this technique the time taken to assemble the best R rectifications from the changes found in the first step appears to be practically independent of the number of inconsistent propositions and so is feasible even for large networks.

The INFERNO output in Fig. 5 shows these processes applied to our continuing example. We look first at proposition B. Since  $t(B)$  was supplied explicitly, B could be made consistent if the given value of  $t(B)$  were reduced to  $1 - f(B)$ , i.e., to 0.625. Alternatively, we could reduce  $f(B)$  to  $1 - t(B)$  or 0.1. But  $f(B)$  was inferred from the given value of  $f(E)$  and the relation "B enables E with strength 0.8" via constraint (1.2). If either  $f(E)$  were 0.28 or the strength of the relation were 0.556,  $f(B)$  would have been 0.1 and B would have been consistent. By this reasoning, there are three changes, any one of which would make the bounds on B consistent:

- reducing  $t(B)$  from 0.9 to 0.625,
- reducing  $f(E)$  from 0.5 to 0.28,
- reducing the strength of "B enables E" from 0.8 to 0.556,

with reluctances 0.275, 0.22, and 0.244, respectively. After changes have been found for all contradictory propositions, INFERNO shows the various rectifications that can be constructed. Notice that reducing  $f(E)$  to 0.28 satisfies all sets of changes and so is a rectification in its own right. This is comforting because the system became inconsistent only when  $f(E)$  was increased beyond 0.28! Again, partially weakening  $f(E)$  to 0.45 will fix proposition C, while reducing the

---

"B" can be resolved by changing:  
Pr("B") from 0.9 to 0.625  
Pr("E") from 0.5 to 0.72  
[ "B" enables "E" with strength 0.8 ] to 0.556

"C" can be resolved by changing:  
Pr("C") from 0.55 to 0.5  
Pr("E") from 0.5 to 0.55

"D" can be resolved by changing:  
Pr("E") from 0.5 to 0.6  
Pr("C") from 0.55 to 0.375  
Pr("B") from 0.9 to 0.688  
[ "B" enables "E" with strength 0.8 ] to 0.611

"E" can be resolved by changing:  
Pr("E") from 0.5 to 0.72  
Pr("B") from 0.9 to 0.625  
[ "B" enables "E" with strength 0.8 ] to 0.556

Alternative rectifications:

Pr("E") from 0.5 to 0.72  
Total reluctance 0.22

[ "B" enables "E" with strength 0.8 ] to 0.556  
Pr("E") from 0.5 to 0.55  
Total reluctance 0.29

[ "B" enables "E" with strength 0.8 ] to 0.556  
Pr("C") from 0.55 to 0.5  
Total reluctance 0.29

Pr("B") from 0.9 to 0.625  
Pr("E") from 0.5 to 0.55  
Total reluctance 0.32

Pr("B") from 0.9 to 0.625  
Pr("C") from 0.55 to 0.5  
Total reluctance 0.32

---

Fig. 5--Changes and rectifications

strength of "B enables E" will fix the other inconsistent propositions, so this pair of changes is also a rectification.

The rectifications constructed by this algorithm are sufficient rather than necessary to make all propositions consistent and so may be excessive when possible interrelationships between the changes are taken into account. In the second rectification above, the reduced strength of 0.556 for the enables relation was computed on the basis of  $f(E)$  being 0.5. Since we also reduced the value of  $f(E)$  to 0.45, the change in strength is slightly more than it need be. Of course, changes can be made one at a time and the residual rectifications recomputed after each change; this would give a new strength of 0.611 for the relation "B enables E".

#### IV. A COMPLETE EXAMPLE

This section returns to the AL/X inference network of Fig. 1 that was discussed in Sec. II, recasts it as INFERNO relations, and runs the network on a case given by Reiter (1981). The purpose of this exercise is to highlight the differences between the Bayesian inferencing employed by AL/X (using the conditional independence assumptions mentioned earlier) and INFERNO's conservative approach with no assumptions.

The comparison of these systems immediately runs afoul of their being based on different world models. The Bayesian approach requires that the prior probability of each proposition A be known and computes the posterior probability of A given all evidence E. Inferential links between propositions A and B are defined in terms of the conditional probabilities  $P(B|A)$  and  $P(B|\sim A)$  that again refer back to the prior probability distributions. In contrast, INFERNO makes no reference to a prior distribution but represents by  $P(A)$  what is known about the proposition A so far. The acquisition of evidence is viewed as a means of further constraining  $P(A)$ . This may appear to be a fine distinction, but consider, for example, the inequality  $P(B|A) \geq X$ . In a Bayesian system, this constrains only the prior distribution, and if  $P(A|E)$  turns out to be 1, it does not follow that  $P(B|E)$  is at least X. INFERNO would interpret the inequality as a constraint that must be satisfied by any assignment of probability intervals in any particular example, and the addition of further evidence could not weaken this constraint. Thus the findings that  $P(A) = 1$  while  $P(B) \leq X$  would be regarded as

inconsistent with the earlier inequality. This difference significantly qualifies the following comparison.

The Boolean relations employed by AL/X are transferable immediately to INFERNO. Let the prior probability of proposition B be  $\text{prior}(B)$ , i.e., the prior odds of B are given by

$$\text{odds}(B) = \text{prior}(B) / (1 - \text{prior}(B))$$

Consider an inferential link from A to B with odds multipliers  $l_s$  (if A is found to be true) and  $l_n$  (if A is found to be false), respectively. If A is found to be true, in the absence of other evidence the posterior probability of B is given by

$$\text{odds}(B|A) = l_s \times \text{odds}(B)$$

and the posterior probability of B is then

$$\text{posterior}(B|A) = \text{odds}(B|A) / (\text{odds}(B|A) + 1)$$

If A is false, we get a corresponding probability by replacing  $l_s$  with  $l_n$  in the odds formula. This link thus becomes a pair of INFERNO relations:

A enables B with strength  $\text{posterior}(B|A)$

B requires A with strength  $(1 - \text{posterior}(B|\sim A))$

Each link is being translated here in isolation, but this should be an accurate translation if the network designer assigned the values of  $l_s$  and  $l_n$  in isolation (as, incidentally, he is advised to do in Reiter (1980, p. 10)) and when A is known with certainty. Differences will

arise when A is known with less than certainty and when multiple inferential links relate to the same proposition, as AL/X then makes use of assumptions that have no counterpart in INFERNO.

The case to be analyzed is defined by probability assignments to the propositions representing raw evidence. AL/X uses quantities called certainty factors (not to be confused with MYCIN's use of the term) in which -5 means false, 5 means true, and 0 means that no information is available. Values in the range 0 to 5 are interpolated linearly between the prior probability of the proposition and 1, with a similar interpolation in the negative range. For the case being analyzed, the certainty factors and probabilities are shown in Table 3. The certainty factor of 0 for proposition RVNOISECOOL presents a minor problem. It could be interpreted as implying the probability shown, but AL/X takes it to mean that the information is either unavailable or is conflicting

Table 3

AL/X CASE FOR ANALYSIS

<u>Proposition</u>	<u>Certainty Factor</u>	<u>Equivalent Probability</u>
RVLIFTIND	5	1
V01CCRCHART	4.5	0.9
RVNOISECOOL	0	0.001
V01SDHP	5	1
PCV302FTBAD	-4	0.0002*
SCIVSHUT	5	1

\*This value is below INFERNO's minimum probability and is treated as 0.



and makes no use of it. This is more accurately reflected in INFERNO as the absence of any statement about the probability, which is the policy followed in this example.

An annotated transcript of the presentation of this case to INFERNO appears as Appendix C. In INFERNO's world view, the information concerning probabilities and relations is inconsistent. PCV302FTBAD is false, for example; thus, so is LIFTPRESS and hence VO1SDHP, directly contradicting another piece of evidence. INFERNO is asked to find alternative rectifications and prints the best ten of them. The first rectification is simply to adjust the probability of PCV302FTBAD to 0.204, and this seems plausible--the information about this proposition was indefinite on the certainty scale and only its very low prior probability of 0.001 caused -4 to be mapped to a near-definite 0.0002 on the probability scale. If this rectification is made, the information becomes consistent, and the findings of AL/X and INFERNO regarding the goal propositions can be compared in Table 4. The results obtained by both systems agree well in general tendency, but AL/X's lie outside the

Table 4

COMPARISON OF FINDINGS BY AL/X AND INFERNO

<u>Proposition</u>	<u>AL/X</u>	<u>INFERNO</u>
SCIVCAUSE	0.909	0.802 - 0.883
PCV302EQERR	0.057	0.204
RVLIFTEARLY	false	0.000 - 0.118
RVSOLSHORT	false	0.059 - 0.199
RVSWSHCRT	false	0.033 - 0.199

probability limits derived by INFERNO; this is not altogether surprising, since one of the input probabilities had to be altered to achieve consistency. INFERNO's bounds also give a good measure of the uncertainty of the various conclusions, as can be noticed from SCIVCAUSE being more tightly bounded than RVSWSHORT.

## V. CONCLUSION

The study reported here was undertaken to develop a useful mechanism for plausible reasoning in the context of uncertain knowledge and culminated in the specification and implementation<sup>7</sup> of INFERNO. A detailed description of the system's attributes was given in Sec. III, but a brief reiteration seems appropriate here. INFERNO is cautious because it does not depend on assumptions about joint probability distributions of propositions, so its conclusions about the probability bounds of propositions are provably correct consequences of the given information. The absence of assumptions would be expected to lead to weaker conclusions, but this tendency is partially offset by enabling sets of propositions that are mutually exclusive or independent to be identified, with the result that probability bounds can be tightened in some cases. The system does not distinguish between hypotheses and evidence and thus can be used for forward (data-driven) inference, backward (hypothesis-driven) inference, or any mixture of these two modes. Finally, INFERNO incorporates a strong notion of the consistency of the information presented to it, and by reasoning backwards about conclusions, it can provide approximate but informative suggestions for remedying any inconsistencies.

As illustrated in the previous section, conventional inference networks can be recast to "equivalent" INFERNO formalisms subject to the

---

<sup>7</sup>INFERNO has been implemented in a mixture of Pascal and C for a VAX 11/780 minicomputer and should be relatively portable among UNIX systems. It is quite economical to run: the entire output in Appendix C, including finding the alternative rectifications, required about 3 seconds of CPU time.

qualification of an underlying non-Bayesian world model. In the oil-rig example, AL/X and INFERNO reached conclusions that were certainly not identical, but which would both support the same diagnosis of the problem. The main advantages of the approach embodied in INFERNO are (1) it requires less information, since it does not need anything corresponding to the prior probability of a hypothesis; (2) it brings out the conflicts that may be implicit in the evidence; and (3) its probability bounds give a measure of the potential error in the conclusions, a feature that has no direct counterpart in Bayesian systems such as AL/X.

INFERNO has also been applied to several more demanding test domains, including a diagnosis network for carburetor malfunctions (also taken from AL/X), several of the published Prospector submodels, transportation planning via a network of unreliable routes, allocating resources among competing but interdependent projects, and assessment of an opponent's poker hand using clues from his bidding and draw. The last case was the only one in which the system proved to be relatively weak. Some prior probabilities in this domain can be calculated, so a Bayesian approach would be potentially more powerful. However, systems such as AL/X do not have any mechanism for enforcing the mutual exclusivity of propositions or for computing accurately the probability of conjunctions or disjunctions of propositions, so the poker domain would probably be a difficult one for them also.

Perhaps the most novel features of INFERNO are its concern for consistency and its rectification-constructing mechanisms. This approach provides a valuable tool to help the user debug the knowledge

or to adapt general rules for a particular problem. In addition it has a somewhat unexpected application in planning whereby the user deliberately introduces contradictions! Consider for example a transportation planning domain in which the goal is to move some combination of loads to their appropriate destinations. One inference from the available information might be that achieving this goal has a probability strictly less than 1. If the goal is also asserted to be satisfied, the information will thus become inconsistent; but the possible rectifications suggested by the system will include combinations of changes sufficient to guarantee achievement of the goal. Thus INFERNO can be used to isolate those characteristics of the domain that bear most significantly on the accomplishment of the planning objective.

In summary, INFERNO is a powerful and flexible tool for dealing with uncertain knowledge, as long as the knowledge can be cast in the form of a fixed set of propositions and relations among them. Nevertheless, several areas for possible improvement suggest themselves. First, it would be handy to have the ability to incorporate Bayesian submodels where the required information is available and the necessary assumptions are found to be valid. Second, real-world tasks often entail budgetary and other numeric but nonprobabilistic constraints, and some formalism is needed for marrying numeric and probability-bounding constraints. Finally, INFERNO is essentially a zeroth-order system in which propositions and relations concern individuals. A quantified relation such as "A(x) enables B(x) for every x" can be represented only as a collection of zeroth-order relations obtained by instantiating over

every individual  $x$ , an unsatisfactory approach if there are many such individuals. We are investigating ways in which an INFERNO-like approach can be moved to a first-order environment.

# Appendix A

## DERIVATION OF PROPAGATION CONSTRAINTS

The following derivations are straightforward manipulations of a few relations, using the identities

$$P(Z) + P(\sim Z) = 1$$

$$t(Z) \leq P(Z) \leq 1 - f(Z)$$

for any proposition Z. Note that the second presumes the consistency of information about Z.

### 1. A enables S with strength X

The interpretation of this relation gives

$$P(S) \geq P(S \& A) = P(A) \times P(S|A) \geq P(A) \times X$$

so

$$P(S) \geq t(A) \times X \dots\dots\dots (1.1)$$

By inverting the earlier relation,

$$P(A) \leq P(S) / X$$

and thus

$$P(\sim A) = 1 - P(A) \geq 1 - (1 - f(S)) / X \dots\dots\dots (1.2)$$

### 2. A negates S

The constraints are immediate consequences of the identities above.

### 3. Conjunction

If A is the conjunction of  $\{S_1, S_2, \dots, S_n\}$ , then  $\sim A$  is the disjunction of  $\{\sim S_1, \sim S_2, \dots, \sim S_n\}$ . Thus the derivations for the various conjunction constraints are mirror images of those for disjunction, interchanging both signs and the bounds t and f throughout.

PRECEDING PAGE BLANK-NOT FILMED

4.1. A disjoins  $\{S_1, S_2, \dots, S_n\}$

The basic relation used here is

$$\max_i P(S_i) \leq P(A) \leq \sum_i P(S_i)$$

The left-hand inequality gives the two constraints

$$P(A) \geq P(S_i) \geq t(S_i) \dots\dots\dots (4.1.1)$$

$$P(\sim S_i) \geq P(\sim A) \geq f(A) \dots\dots\dots (4.1.4)$$

The right-hand inequality gives

$$P(\sim A) \geq 1 - \sum_i (1 - f(S_i)) \dots\dots\dots (4.1.2)$$

and by rewriting it in the form

$$P(S_i) \geq P(A) - \sum_{j \neq i} P(S_j)$$

we get the constraint

$$P(S_i) \geq t(A) - \sum_{j \neq i} (1 - f(S_j)) \dots\dots\dots (4.1.3)$$

4.2. A disjoins-independent  $\{S_1, S_2, \dots, S_n\}$

The rule for combining the probabilities of independent events is

$$P(\sim A) = \prod_i P(\sim S_i)$$

giving

$$P(A) = 1 - P(\sim A) \geq 1 - \prod_i (1 - t(S_i)) \dots\dots\dots (4.2.1)$$

$$P(\sim A) \geq \prod_i f(S_i) \dots\dots\dots (4.2.2)$$

Rewriting the rule above as

$$P(\sim S_i) = P(\sim A) / \prod_{j \neq i} P(\sim S_j)$$

gives

$$P(\sim S_i) \geq f(A) / \prod_{j \neq i} (1 - t(S_j)) \dots\dots\dots (4.2.4)$$

$$P(S_i) = 1 - P(\sim S_i) \geq 1 - (1 - t(A)) / \prod_{j \neq i} f(S_j) \dots\dots (4.2.3)$$



4.3. A disjoints-exclusive  $\{S_1, S_2, \dots, S_n\}$

The rule in this case ..

$$P(A) = \sum_i P(S_i)$$

hence

$$P(A) \geq \sum_i t(S_i) \dots\dots\dots (4.3.1)$$

$$P(\sim A) = 1 - P(A) \geq 1 - \sum_i (1 - f(S_i)) \dots\dots\dots (4.3.2)$$

Rearranging the rule,

$$P(S_i) = P(A) - \sum_{j \neq i} P(S_j)$$

which gives

$$P(S_i) \geq t(A) - \sum_{j \neq i} (1 - f(S_j)) \dots\dots\dots (4.3.3)$$

$$P(\sim S_i) = 1 - P(S_i) \geq f(A) + \sum_{j \neq i} t(S_j) \dots\dots\dots (4.3.4)$$

5.  $\{S_1, S_2, \dots, S_n\}$  mutually exclusive

The rule is

$$\sum_i P(S_i) \leq 1$$

or

$$P(S_i) \leq 1 - \sum_{j \neq i} P(S_j)$$

Thus

$$P(\sim S_i) = 1 - P(S_i) \geq \sum_{j \neq i} t(S_j) \dots\dots\dots (5.1)$$

Appendix B

CONSTRAINTS FOR BACKING UP INCONSISTENCIES

A enables S with strength X:

$$t(A) \leq t(S) / X$$

$$f(S) \leq 1 - (1 - f(A)) \times X$$

A negates S:

$$t(A) = f(S)$$

$$f(A) = t(S)$$

A disjoins  $\{S_1, S_2, \dots, S_n\}$ :

$$t(A) \leq t(S_i) + \sum_{j \neq i} (1 - f(S_j))$$

$$f(A) \leq f(S_i)$$

$$t(S_i) \leq t(A)$$

$$f(S_i) \leq t(S_j) + \sum_{k \neq i, j} (1 - f(S_k)) + 1 - t(A)$$

$$f(S_i) \leq f(A) + \sum_{j \neq i} (1 - f(S_j))$$

A disjoins-independent  $\{S_1, S_2, \dots, S_n\}$

$$t(A) \leq 1 - (1 - t(S_i)) \times \prod_{j \neq i} f(S_j)$$

$$f(A) \leq f(S_i) \times \prod_{j \neq i} (1 - t(S_j))$$

$$t(S_i) \leq 1 - (1 - t(A)) / \prod_{j \neq i} (1 - t(S_j))$$

$$t(S_i) \leq 1 - f(A) / (f(S_j) \times \prod_{k \neq i, j} (1 - t(S_k)))$$

$$f(S_i) \leq f(A) / \prod_{j \neq i} f(S_j)$$

$$f(S_i) \leq (1 - t(A)) / ((1 - t(S_j)) \times \prod_{k \neq i, j} f(S_k))$$

A disjoins-exclusive  $\{S_1, S_2, \dots, S_n\}$

$$t(A) \leq t(S_i) + \sum_{j \neq i} (1 - f(S_j))$$

$$f(A) \leq f(S_i) - \sum_{j \neq i} t(S_j)$$

$$t(S_i) \leq t(A) - \sum_{j \neq i} t(S_j)$$

$$t(S_i) \leq f(S_j) - \sum_{k \neq i, j} t(S_k) - f(A)$$

$$f(S_i) \leq f(A) + \sum_{j \neq i} (1 - f(S_j))$$

$$f(S_i) \leq t(S_j) + \sum_{k \neq i, j} (1 - f(S_k)) + 1 - t(A)$$

$\{S_1, S_2, \dots, S_n\}$  mutually exclusive

$$t(S_i) \leq f(S_j) - \sum_{k \neq i, j} t(S_k)$$

Appendix C

INFERNO RUN WITH AL/X EXAMPLE

```
+-----+
| Define the various relations |
| corresponding to AL/X links. |
+-----+
```

"VOISDHP" enables "LIFTPRESS" with strength 0.668;  
"LIFTPRESS" requires "VOISDHP" with strength 1.000;

"VOICCRCHART" enables "LIFTPRESS" with strength 0.980;  
"LIFTPRESS" requires "VOICCRCHART" with strength 1.000;

"LP&SCIVSHUT" conjoins {"LIFTPRESS", "SCIVSHUT"};

"LP&SCIVSHUT" enables "SCIVCAUSE" with strength 0.909;  
"SCIVCAUSE" requires "LP&SCIVSHUT" with strength 1.000;

"PCV302FTBAD" enables "PCV302EQERR" with strength 0.968;  
"PCV302EQERR" requires "PCV302FTBAD" with strength 1.000;

"LIFTPRESS" enables "PCV302EQERR" with strength 0.231;  
"PCV302EQERR" requires "LIFTPRESS" with strength 0.998;

"LIFTPRESS" enables "RELVLIFT" with strength 0.668;  
"RELVLIFT" requires "LIFTPRESS" with strength 1.000;

"RVNOISECOOL" enables "RELVLIFT" with strength 0.501;  
"RELVLIFT" requires "RVNOISECOOL" with strength 0.997;

"RVLIFTIND" enables "RELVLIFT" with strength 0.801;  
"RELVLIFT" requires "RVLIFTIND" with strength 1.000;

"NLIFTPRESS" negates "LIFTPRESS";

"LIFT&NLIFTP" conjoins {"RELVLIFT", "NLIFTPRESS"};

"LIFT&NLIFTP" enables "RVLIFTEARLY" with strength 0.667;  
"RVLIFTEARLY" requires "LIFT&NLIFTP" with strength 1.000;

"NRELVLIFT" negates "RELVLIFT";

"NLIFT&IND" conjoins {"RVLIFTIND", "NRELVLIFT"};

"NLIFT&IND" enables "RVSOLSHORT" with strength 0.500;  
"RVSOLSHORT" requires "NLIFT&IND" with strength 1.000;

"NLIFT&IND" enables "RVSWSHORT" with strength 0.286;  
"RVSWSHORT" requires "NLIFT&IND" with strength 1.000;

```
+-----+
| Enter the probabilities |
| defining this particular case. |
+-----+
```

```
assume "RVLIFTIND";
assume "VO1CCRCHART" with probability 0.9;
assume "VO1SDHP";
assume "PCV302FTBAD" false;
assume "SCIVSHUT";

show events;

** "VO1SDHP": range 1 - 0
    LB by assumption
    UB from "LIFTPRESS" being false, by enables-inverse

** "LIFTPRESS": range 0.882 - 0
    LB from Pr("VO1CCRCHART") >= 0.9, by enables
    UB from "PCV302EQERR" being false, by enables-inverse

** "VO1CCRCHART": range 0.9 - 0
    LB by assumption
    UB from "LIFTPRESS" being false, by enables-inverse

"LP&SCIVSHUT": false
    UB from "LIFTPRESS" being false, by conjoins

"SCIVSHUT": true
    LB by assumption

"SCIVCAUSE": false
    UB from "LP&SCIVSHUT" being false, by requires

** "PCV302FTBAD": range 0.204 - 0
    LB from Pr("PCV302EQERR") >= 0.204, by requires-inverse
    UB by assumption

** "PCV302EQERR": range 0.204 - 0
    LB from Pr("LIFTPRESS") >= 0.882, by enables
    UB from "PCV302FTBAD" being false, by requires

** "RELVLIFT": range 0.801 - 0
    LB from "RVLIFTIND" being true, by enables
    UB from "LIFTPRESS" being false, by requires

** "RVNOISECOOL": range 0.8 - 0
    LB from Pr("RELVLIFT") >= 0.801, by requires-inverse
    UB from "RELVLIFT" being false, by enables-inverse
```

\*\* "RVLIFTIND": range 1 - 0  
LB by assumption  
UB from "RELVLIFT" being false, by enables-inverse

\*\* "NLIFTPRESS": range 1 - 0.118  
LB from "LIFTPRESS" being false, by negates  
UB from Pr("LIFTPRESS") >= 0.882, by negates

"LIFT&NLIFTP": false  
UB from "RELVLIFT" being false, by conjoins

"RVLIFTEARLY": false  
UB from "LIFT&NLIFTP" being false, by requires

\*\* "NRELVLIFT": range 1 - 0.199  
LB from "RELVLIFT" being false, by negates  
UB from Pr("RELVLIFT") >= 0.801, by negates

\*\* "NLIFT&IND": range 0.1 - 0  
LB from "NRELVLIFT" being true, by conjoins  
UB from "RVLIFTIND" being false, by conjoins

\*\* "RVSOLSHORT": range 0.05 - 0  
LB from Pr("NLIFT&IND") >= 0.1, by enables  
UB from "NLIFT&IND" being false, by requires

\*\* "RVSWSHORT": range 0.029 - 0  
LB from Pr("NLIFT&IND") >= 0.1, by enables  
UB from "NLIFT&IND" being false, by requires

```
+-----+  
| Find ways of making the |  
| information consistent. |  
+-----+
```

show rectifications;

"VO1SDHP" can be resolved by changing:

Pr("VO1SDHP") from 1 to 0  
[ "VO1SDHP" enables "LIFTPRESS" with strength 0.668 ] to 0  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.154

"LIFTPRESS" can be resolved by changing:

[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("VO1CCRCHART") from 0.9 to 0  
[ "VO1CCRCHART" enables "LIFTPRESS" with strength 0.98 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.204

"VO1CCRCHART" can be resolved by changing:

Pr("VO1CCRCHART") from 0.9 to 0  
[ "VO1CCRCHART" enables "LIFTPRESS" with strength 0.98 ] to 0  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.204

"PCV302FTBAD" can be resolved by changing:

Pr("PCV302FTBAD") from 0 to 0.204  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("VO1CCRCHART") from 0.9 to 0  
[ "VO1CCRCHART" enables "LIFTPRESS" with strength 0.98 ] to 0

"PCV302EQERR" can be resolved by changing:

Pr("PCV302FTBAD") from 0 to 0.204  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("VO1CCRCHART") from 0.9 to 0  
[ "VO1CCRCHART" enables "LIFTPRESS" with strength 0.98 ] to 0

"RELVLIFT" can be resolved by changing:

Pr("RVLIFTIND") from 1 to 0  
[ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.185

"RVNOISECOOL" can be resolved by changing:

[ "RVNOISECOOL" enables "RELVLIFT" with strength 0.501 ] to 0  
[ "RELVLIFT" requires "RVNOISECOOL" with strength 0.997 ] to 0.199  
Pr("RVLIFTIND") from 1 to 0.004  
[ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0.003  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.093

"RVLIFTIND" can be resolved by changing:

Pr("RVLIFTIND") from 1 to 0  
[ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.185

"NLIFTPRESS" can be resolved by changing:

[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("VO1CCRCHART") from 0.9 to 0  
[ "VO1CCRCHART" enables "LIFTPRESS" with strength 0.98 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.204

"NRELVLIFT" can be resolved by changing:

Pr("RVLIFTIND") from 1 to 0  
[ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.185

"NLIFT&IND" can be resolved by changing:

Pr("RVLIFTIND") from 1 to 0  
[ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.018

"RVSOLSHORT" can be resolved by changing:

[ "NLIFT&IND" enables "RVSOLSHORT" with strength 0.5 ] to 0  
Pr("RVLIFTIND") from 1 to 0  
[ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.009

"RVSWSHORT" can be resolved by changing:

[ "NLIFT&IND" enables "RVSWSHORT" with strength 0.286 ] to 0  
Pr("RVLIFTIND") from 1 to 0  
[ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Pr("PCV302FTBAD") from 0 to 0.005

Alternative rectifications:

Pr("PCV302FTBAD") from 0 to 0.204  
Total reluctance 0.20

[ "LIFTPRESS" enables "PCV302EQERR" with strength 0.231 ] to 0  
Total reluctance 0.23

Pr("PCV302FTBAD") from 0 to 0.185  
Pr("V01CCRCHART") from 0.9 to 0  
Total reluctance 1.09

Pr("PCV302FTBAD") from 0 to 0.185  
[ "V01CCRCHART" enables "LIFTPRESS" with strength 0.98 ] to 0  
Total reluctance 1.17

Pr("PCV302FTBAD") from 0 to 0.154  
Pr("V01CCRCHART") from 0.9 to 0  
[ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
Total reluctance 1.86

Pr("PCV302FTBAD") from 0 to 0.154  
[ "V01CCRCHART" enables "LIFTPRESS" with strength 0.98 ] to 0  
[ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
Total reluctance 1.94

Pr("PCV302FTBAD") from 0 to 0.154  
Pr("V01CCRCHART") from 0.9 to 0  
Pr("RVLIFTIND") from 1 to 0  
Total reluctance 2.05



Pr("PCV302FTBAD") from 0 to 0.154  
 [ "V01CCRCHART" enables "LIFTPRESS" with strength 0.98 ] to 0  
 Pr("RVLIFTIND") from 1 to 0  
 Total reluctance 2.13

[ "V01SDHP" enables "LIFTPRESS" with strength 0.668 ] to 0  
 Pr("V01CCRCHART") from 0.9 to 0  
 [ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
 Total reluctance 2.37

[ "V01SDHP" enables "LIFTPRESS" with strength 0.668 ] to 0  
 [ "V01CCRCHART" enables "LIFTPRESS" with strength 0.98 ] to 0  
 [ "RVLIFTIND" enables "RELVLIFT" with strength 0.801 ] to 0  
 Total reluctance 2.45

```
+-----+
| Adopt the first rectification |
| given above.                  |
+-----+
```

assume "PCV302FTBAD" with probability 0.204;

show events;

"V01SDHP": true  
 LB by assumption

"LIFTPRESS": range 0.882 - 0.883  
 LB from Pr("V01CCRCHART")  $\geq$  0.9, by enables  
 UB from Pr("PCV302EQERR")  $\leq$  0.204, by enables-inverse

"V01CCRCHART": range 0.9 - 0.901  
 LB by assumption  
 UB from Pr("LIFTPRESS")  $\leq$  0.883, by enables-inverse

"LP&SCIVSHUT": range 0.882 - 0.883  
 LB from Pr("LIFTPRESS")  $\geq$  0.882, by conjoins  
 UB from Pr("LIFTPRESS")  $\leq$  0.883, by conjoins

"SCIVSHUT": true  
 LB by assumption

"SCIVCAUSE": range 0.802 - 0.883  
 LB from Pr("LP&SCIVSHUT")  $\geq$  0.882, by enables  
 UB from Pr("LP&SCIVSHUT")  $\leq$  0.883, by requires

"PCV302FTBAD": range 0.204 - 0.204  
 LB by assumption  
 UB by assumption

"PCV302EQERR": range 0.204 - 0.204  
LB from Pr("LIFTPRESS")  $\geq$  0.882, by enables  
UB from Pr("PCV302FTBAD")  $\leq$  0.204, by requires

"RELVLIFT": range 0.801 - 0.883  
LB from "RVLIFTIND" being true, by enables  
UB from Pr("LIFTPRESS")  $\leq$  0.883, by requires

"RVNOISECOOL": range 0.8 - 1  
LB from Pr("RELVLIFT")  $\geq$  0.801, by requires-inverse

"RVLIFTIND": true  
LB by assumption

"NLIFTPRESS": range 0.117 - 0.118  
LB from Pr("LIFTPRESS")  $\leq$  0.883, by negates  
UB from Pr("LIFTPRESS")  $\geq$  0.882, by negates

"LIFT&NLIFTP": range 0 - 0.118  
UB from Pr("NLIFTPRESS")  $\leq$  0.118, by conjoins

"RVLIFTEARLY": range 0 - 0.118  
UB from Pr("LIFT&NLIFTP")  $\leq$  0.118, by requires

"NRELVLIFT": range 0.117 - 0.199  
LB from Pr("RELVLIFT")  $\leq$  0.883, by negates  
UB from Pr("RELVLIFT")  $\geq$  0.801, by negates

"NLIFT&IND": range 0.117 - 0.199  
LB from Pr("NRELVLIFT")  $\geq$  0.117, by conjoins  
UB from Pr("NRELVLIFT")  $\leq$  0.199, by conjoins

"RVSOLSHORT": range 0.059 - 0.199  
LB from Pr("NLIFT&IND")  $\geq$  0.117, by enables  
UB from Pr("NLIFT&IND")  $\leq$  0.199, by requires

"RVSWSHORT": range 0.033 - 0.199  
LB from Pr("NLIFT&IND")  $\geq$  0.117, by enables  
UB from Pr("NLIFT&IND")  $\leq$  0.199, by requires

# REFERENCES

1. Barnett, J.A., "Computational Methods for a Mathematical Theory of Evidence," Proc. 7th International Joint Conf. Artificial Intelligence, Vancouver, 1981, pp. 868-875.
2. Duda, R.O., P.E. Hart, and Nils Nilsson, Subjective Bayesian Methods for Rule-Based Inference Systems, Technical Note 124, Artificial Intelligence Center, SRI International, 1976.
3. Duda, R.O., P.E. Hart, N.J. Nilsson, R. Reboh, J. Slocum, and G.L. Sutherland, Development of a Computer-Based Consultant for Mineral Exploration, SRI International, 1977.
4. Duda, R.O., J. Gaschnig, and P.E. Hart, "Model Design in the Prospector Consultant System for Mineral Exploration," Expert Systems in the Micro Electronic Age (D. Michie, ed.), Edinburgh University Press, 1979.
5. Duda, R.O., P.E. Hart, K. Konolige, and R. Reboh, A Computer-Based Consultant for Mineral Exploration, SRI International, 1979.
6. Feigenbaum, E.A., "Themes and Case Studies of Knowledge Engineering," Expert Systems in the Micro Electronic Age (D. Michie, ed.), Edinburgh University Press, 1979.
7. Friedman, L., "Extended Plausible Inference," Proc. 7th International Joint Conf. Artificial Intelligence, Vancouver, 1981, pp. 487-495.
8. Garvey, T.D., J.D. Lowrance, and M.A. Fischler, "An Inference Technique for Integrating Knowledge from Disparate Sources," Proc. 7th International Joint Conf. Artificial Intelligence, Vancouver, 1981, pp. 319-325.
9. Gaschnig, J., Development of Uranium Exploration Models for the Prospector Consultant System, SRI International, 1980.
10. Gaschnig, J., "Prospector: An Expert System for Mineral Exploration," State of the Art Report on Machine Intelligence (A. Bond, ed.), London: Pergamon Infotech, 1981.
11. Hayes-Roth, F., "Probabilistic Dependencies in a System for Truth Maintenance and Belief Revision," unpublished working paper, The Rand Corporation.
12. Hayes-Roth, F., D.A. Waterman, and D.B. Lenat, "Principles of Pattern-Directed Inference Systems", Pattern-Directed Inference Systems (D.A. Waterman and F. Hayes-Roth, eds.), New York: Academic Press, 1978.

13. Ishizuka, M., and J.T.P. Yao, "Inexact Inference for Rule-Based Damage Assessment of Existing Structures," Proc. 7th International Joint Conf. Artificial Intelligence, Vancouver, 1981, pp. 837-842.
14. Konolige, K., "An Information-Theoretic Approach to Subjective Bayesian Inference in Rule-Based Systems," unpublished draft, SRI International.
15. Paterson, A., AL/X User Manual, Intelligent Terminals Ltd., Oxford, 1981.
16. Pearl, J., Distributed Bayesian Processing for Belief Maintenance in Hierarchical Inference Systems, Cognitive Systems Laboratory, UCLA, Report UCLA-ENG-CSL-82-11, 1982.
17. Pednault, E.D.P., S.W. Zucker, and L.V. Muresan, "On the Independence Assumption Underlying Subjective Bayesian Inference," Artificial Intelligence, Vol. 16, 1981, pp. 213-222.
18. Reiter, J., AL/X: An Expert System using Plausible Inference, Intelligent Terminals Ltd., Oxford, 1980.
19. Reiter, J., "AL/X: An Inference System for Probabilistic Reasoning," M.Sc. Thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, 1981.
20. Robinson, J.A., "A Machine Oriented Logic Based on the Resolution Principle," Journal of the Association for Computing Machinery, Vol. 12, 1965, pp. 23-41.
21. Shortliffe, E.H., and B.G. Buchanan, "A Model of Inexact Reasoning in Medicine," Mathematical Biosciences, Vol. 23, 1975, pp. 351-379.
22. Stallman, R.M., and G.J. Sussman, "Forward Reasoning and Dependency-Directed Backtracking in a System for Computer-Aided Circuit Analysis," Artificial Intelligence, Vol. 9, 1977, pp. 135-196.
23. Szolovits, P.S., and S.G. Pauker, "Categorical and Probabilistic Reasoning in Medical Diagnosis," Artificial Intelligence, Vol. 11, 1978, pp. 115-144.
24. van Melle, W., "A Domain-Independent Production Rule System for Consultation Programs," Proc. 6th International Joint Conf. Artificial Intelligence, Tokyo, 1979.
25. van Melle, W., A Domain-Independent System that Aids in Constructing Knowledge-Based Consultation Programs, Department of Computer Science, Stanford University, Report STAN-CS-80-820, 1980.
26. Zadeh, L.A., "A Theory of Approximate Reasoning," Machine Intelligence, Vol. 9, (J.E. Hayes, D. Michie, and L.I. Mikulich, eds.), London: Ellis Horwood, 1979.